Automating Mid- and Long-Range Scheduling for the NASA Deep Space Network

Mark D. Johnston* and Daniel Tran*

*Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena CA USA 91109
mark.d.johnston & daniel.tran @jpl.nasa.gov

Abstract

NASA has recently deployed a new mid-range scheduling system for the antennas of the Deep Space Network (DSN), called Service Scheduling Software, or S^3. This system was designed and deployed as a modern web application containing a central scheduling database integrated with a collaborative environment, exploiting the same technologies as social web applications but applied to a space operations context. This is highly relevant to the DSN domain since the network schedule of operations is developed in a peer-to-peer negotiation process among all users of the DSN. These users represent not only NASA's deep space missions, but also international partners and ground-based science and calibration users. The initial implementation of S^3 is complete and the system has been operational since July 2011. This paper describes some key aspects of the S^3 system and on the challenges of modeling complex scheduling requirements and the ongoing extension of S^3 to encompass long-range planning, downtime analysis, and forecasting, as the next step in developing a single integrated DSN scheduling tool suite to cover all time ranges.

1. Introduction

The NASA Deep Space Network (DSN) consists of three large complexes of antennas, spaced roughly evenly in longitude around the world at Goldstone, California; Madrid, Spain; and Canberra, Australia. Each complex contains one 70 meter antenna along with a number of 34 meter and smaller antennas, as well as the electronics and networking infrastructure to command and control the antennas and to communicate with various mission control centers. Table 1 summarizes the DSN Deep Space Communications Complexes (DSCC) including their locations, antennas, and supported communications bands; for more extensive background on the DSN, refer to [1, 2].

All NASA planetary and deep space missions, as well as many international missions, communicate to Earth through the DSN. In some cases, missions closer to Earth also use the DSN, some routinely, others on an occasional basis. The capabilities of the DSN make it a scientific facility in its own right, so it is used for radio astronomy (including very long baseline interferometry) as well as radio science investigations. At present, there are 37 regular distinct users of DSN, who together schedule about 500 activities per week. Over the next few decades, utilization of the DSN is expected to grow significantly, with more missions operating, higher data rates and link complexities, and the possibility of manned mission support. In addition, there is significant pressure to reduce ongoing costs while maintaining an around-the-clock operational capability.

In this paper we first give a general overview of the DSN and the nature of its scheduling problem, followed by a brief description of the scheduling process and software systems (Section 2). We then describe some recent extensions to the Service Scheduling Software (SSS, or S^3), covering both additional request types (Section 3) and long-range planning and forecasting (Section 4). Finally we summarize progress to date and plans for future work in our conclusions (Section 5).

2. DSN Scheduling: Process and Software

The DSN scheduling process consists of three phases, which do not have sharply defined boundaries. In this section we briefly describe these phases as they exist today; later in this paper we discuss plans for how they may change in the future.

- *Long-Range Planning and Forecasting*. In today's system, long-range planning is based on user-provided high-level requirements, specified in the form of a spreadsheet that is interpreted by analysts and entered into a database at JPL. The forecast software employs a statistical allocation method[3, 4] to estimate when these requirements translate into DSN loading over

Complex	GDSCC	CDSCC	MDSCC
Location	Goldstone, California, USA	Canberra, Australia	Madrid, Spain
Longitude	117° W	149° E	4° W
Latitude	35° N	35° S	40° N
Antennas	1 - 70m 5 - 34m	1 - 70m 2 - 34m	1 - 70m 3 - 34m
Capabilities	S, X, Ka	S, X Ka (D/L) only	S, X Ka (D/L) only

Table 1. Deep Space Network (DSN) communications complexes and some of their characteristics.

various time frames. Long-range planning has several major purposes:

  - studies and analyses: periods of particular interest or concern are examined to determine where there is likely contention among missions, for example around launches or critical mission events (maneuvers, planetary orbit insertion or landings), or when construction of a new DSN antenna is under investigation
  - downtime analysis: identifying periods of time when necessary antenna or other maintenance can be scheduled, attempting to minimize the impact on missions
  - future mission analysis: in proposal phase, missions can request analysis of their proposed DSN coverage as part of assessing and costing proposals for new missions

  The time range for long-range planning is generally six months or more into the future, sometimes as much as years.

- *Mid-Range Scheduling*. The mid-range scheduling phase is when detailed user requirements are specified, integrated, negotiated, and all tracking activities finalized in the schedule. Starting at roughly 4-5 months before execution, users specify their detailed scheduling requirements on a rolling weekly basis. These requirements include:
  - tracking time and services required
  - constraining time intervals and relationships
  - visibility constraints
  - flexibilities

  More details on these various types of scheduling requirements are provided elsewhere[5-7]. Once the deadline passes and all requirements are in, the full set is integrated into an initial schedule in which conflicts are reduced by taking advantage of whatever flexibilities have been specified. There follows an optimization step where an experienced DSN scheduler interactively edits the schedule and further reduces conflicts by taking advantage of unspecified flexibilities and making further adjustments. At the conclusion of this phase, the schedule usually contains a fewer than 30 conflicting sets of activities. It is then released to the scheduling user community who negotiate to reduce conflicts and further optimize coverage for their missions. This phase generally lasts 7-8 working days, after which the schedule is conflict free or has only waived conflicts for specific reasons. This is considered the "negotiated schedule" that missions use to plan their integrated ground and spacecraft activities, including the development of onboard command loads based in part on the DSN schedule. Following this point, changes to the schedule may still occur, but new conflicts may not be introduced. There is a continuing low level of no-impact changes and negotiated changes that occur all the way down to real time.

- *Near Real-time Scheduling*. The near real-time phase of DSN scheduling starts roughly eight weeks from execution and includes the period through execution of all the scheduled activities. Late changes may occur for various reasons (sometimes impacting the mid-range phase as well):
  - users may have additional information or late changes to requirements for a variety of reasons
  - DSN assets (antennas, equipment) may experience unexpected downtimes that require adjustments to the schedule to accommodate
  - spacecraft emergencies may occur that require extra tracking or changes to existing scheduled activities

  For many missions that are sequenced well in advance, late changes cannot be readily accommodated.

The DSN scheduling software systems represent a collection built over many years and interfaced in a very heterogeneous manner. At the present time, the different stages of the scheduling process are mostly supported by different sets of tools and databases. The DSN has undertaken an overall unification and simplification of the scheduling software systems[5, 6, 8, 9], of which the first part has been operational since mid-2011. This is called the *Service Scheduling Software* (SSS, or S^3) and has initially been applied to the mid-range phase of the process described above.

S^3 provides support for all the key elements of the mid-range process, based on a Javascript-based HTML5 web application and integrated database[10] (see Figure 1). Users can directly enter their own scheduling requirements and verify their correctness before the submission deadline. The database in which requirements are stored is logically divided into "master" and "workspace" areas. There is a single master schedule representing mission-approved requirements and DSN activities (tracks). Each user can create an arbitrary number of workspace schedules, initially either empty or based on the contents of the master, within which they can conduct studies and 'what if' investigations, or keep a baseline for comparison with the master. These workspaces are by default private to the individual user, but can be shared as readable or read-write to any number of other users. Shared workspaces can be viewed and updated in realtime: while there can only be one writer at a time, any number of other users can view a workspace and see it automatically update as changes are made. These aspects of the web application architecture and database design support the collaborative and shared development nature of the DSN schedule.

In addition, S^3 offers specialized features to facilitate collaboration, including an integrated wiki for annotated discussion of negotiation proposals, integrated chat, notifications of various events, and a propose/concur/reject/counter workflow manager to support change proposals. Details on the design and use of the S^3 collaboration features are provided elsewhere[10].

Underlying the web application and database is a scheduling automation component, the DSN Scheduling Engine[11] (DSE). The DSE provides a range of functions

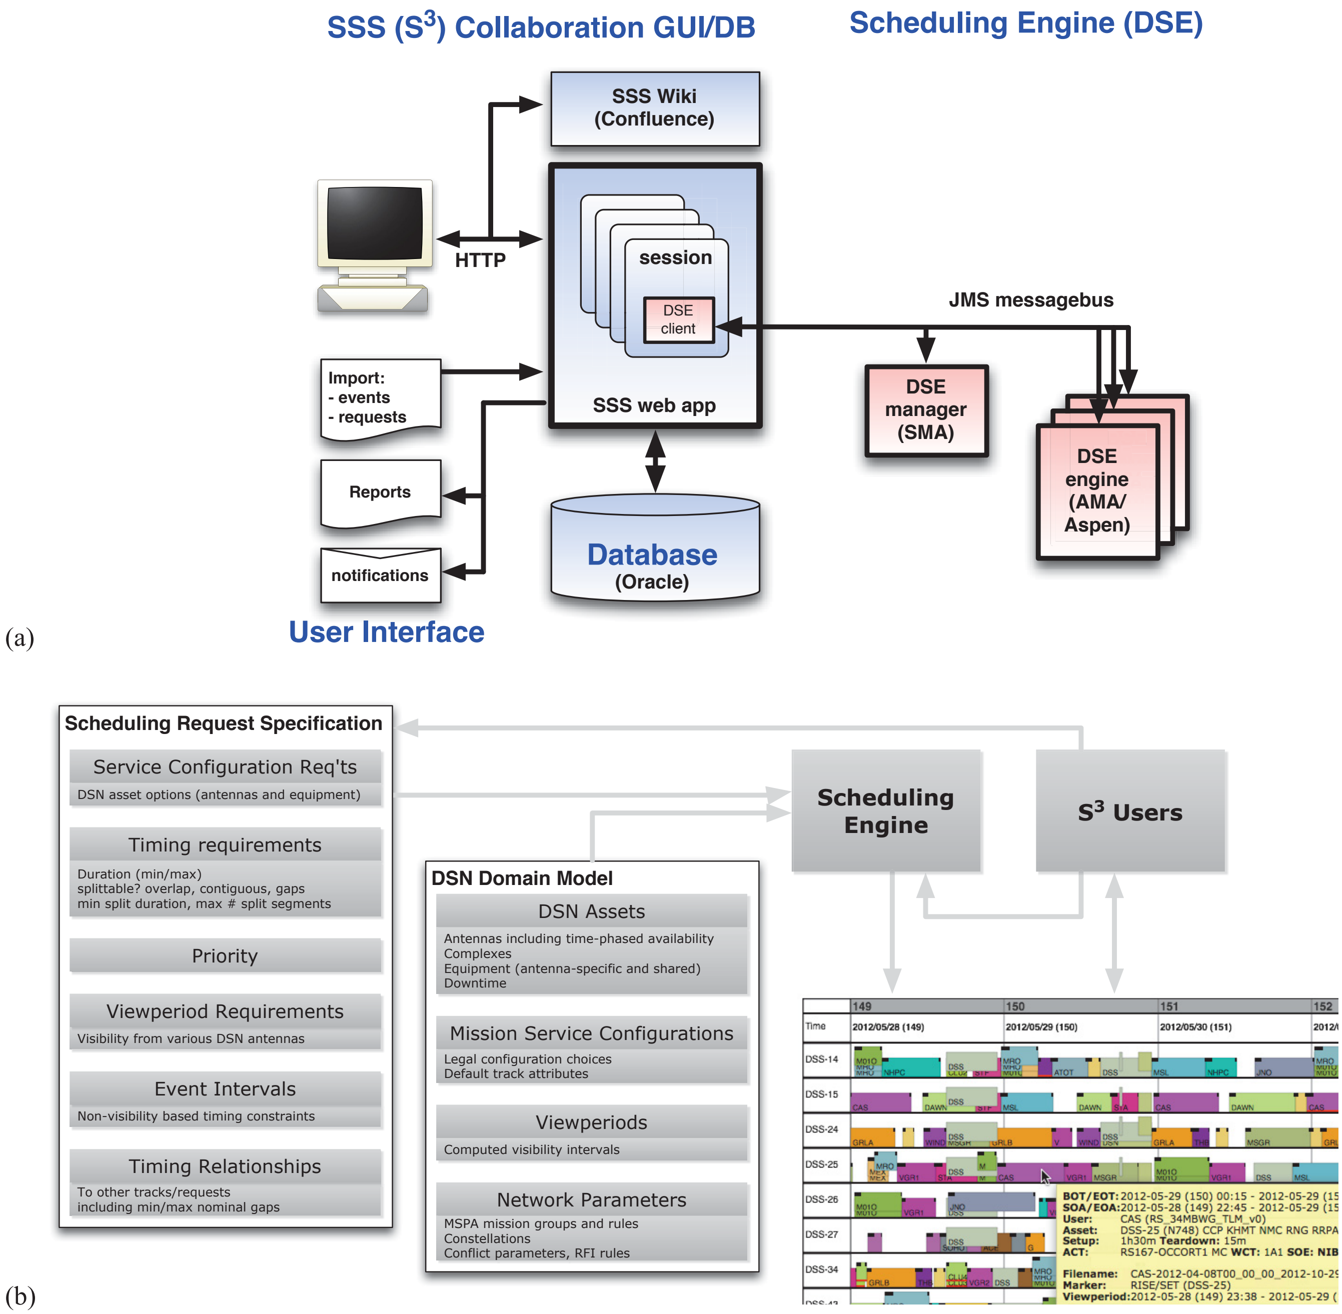

Figure 1: (a) Block diagram of S^3 software elements showing how the user interacts with the system; (b) major elements of the S^3 request specification and domain model, also showing the S^3 HTML5 canvas GUI

based on the semantics of the DSN scheduling domain, including:

- expanding scheduling requests and requirements into tracking or other activities
- checking for and identifying conflicts in the schedule, i.e. situations that violate any DSN scheduling rules
- checking for and identifying requirement violations in the schedule, i.e. situations where activities in the schedule do not meet the user's specified requirements and constraints
- deconfliction algorithms that attempt to reduce conflicts or violations while preserving satisfied requirements

The DSE is based on a distributed session-oriented infrastructure running the ASPEN planning system[12] with a DSN domain adaptation layer.

3. Extended Scheduling Request Types

The initial deployment of S^3 has focussed on the most frequently encountered types of scheduling request types, which *directly* affect how DSN antenna allocations are to be constructed. Direct requests and requirements specify such attributes as:

- tracking duration, and duration flexibility
- whether activities can be split, and, if so, whether the split segments must be overlapping, contiguous, or separated by gaps
- which antennas and equipment combinations may be used to satisfy the requirement
- timing linkages among activities
- constraints on when activities can be scheduled based on occurrence of specified events

Since the initial deployment of S^3, work has been ongoing on a second category of scheduling requirement, which *indirectly* affect allocations in a non-local manner. By this we mean that an extended time period and multiple activities may have to be examined to determine whether some preferred condition is satisfied. These conditions can have a varying degree of preference, ranging from very high to quite weak. It can also be the case that there is a tradeoff between satisfying these types of requirements *vs*. the direct requirements noted above. Examples of these indirect requirements follow:

- 3 out of every 10 tracking passes must be scheduled at the Canberra complex (i.e. in the southern hemisphere)
- there must be scheduled 6 hours of uplink per day and 12 hours of downlink per day, no matter how divided among different antennas, measured midnight-to-midnight UTC
- there must be at least 24 hours of tracking time scheduled per week, added up over four related missions
- downlink tracks of sufficient duration must be scheduled to ensure that onboard recorder capacity is not exceeded

We have denoted these types of scheduling requests as *timeline* constraints or preferences, since they are best assessed by considering the overall timeline of activities (or subset of activities) for a DSN service user over some time period. Table 2 includes a more detailed list of major timeline requirement types and their parameters.

Because these requests have a varying degree of preference, and therefore need to be accessible to the judgement of the scheduling users, we have pursued their incorporation into S^3 in two phases:

- as integrated with the scheduling system graphical user interface (GUI), for visualization along with the actual schedule itself
- as incorporated into the DSE algorithm set, for invocation as strategies or heuristic repair & rescheduling options that can be included or not into the normal scheduling process

Request Type	Examples	Parameters
Total Time	• 8 hours of tracking per day • 6 hours of uplink tracking each midnight to midnight UTC • 24 hours of specific activity types per week summed over four different but related spacecraft	• mission(s) • service aliases • time frame (1 day, 1 week, etc.) • min/max tracking times with yellow/red limits
Tracking Gaps	• 6-12 hour gap between tracks, measured mid-point to mid-point • gaps no greater than 8 hours measured EOT to BOT	• mission • service aliases • min track gap • max track gap • yellow limits • measured by (BOT-BOT, EOT-EOT, mid track to mid track)
DSN Complex Distribution	• 3 of 10 tracks per week must be scheduled at Canberra DSN complex • at least one track per week must be scheduled at each DSN complex	• mission • duration • list of (complex, count)
Recorder	• do not exceed onboard recorder volume capacity limit	• mission • track overhead duration • recorder collection rate (X units/s) • yellow/red recorder max capacity • recorder downlink rates (antenna, downlink rate X units/s) • initialization rule

Table 2. Timeline requirement types, with examples and parameters.

Integration with the S^3 GUI has built upon the newly deployed S^3 HTML5 canvas-based GUI (see Figure 1b [7]), which has enabled the rapid extension of the GUI to additional visualization elements. We provide examples of the visualization of each of the major categories of timeline requirements below.

The Total Time timeline requirement applies to about 25% of the DSN user set, but over a wide range of timescales, from a full week on down to a fraction of a single day. An example for the GRAIL A/B mission (two spacecraft in lunar orbit) is shown in Figure 2a.

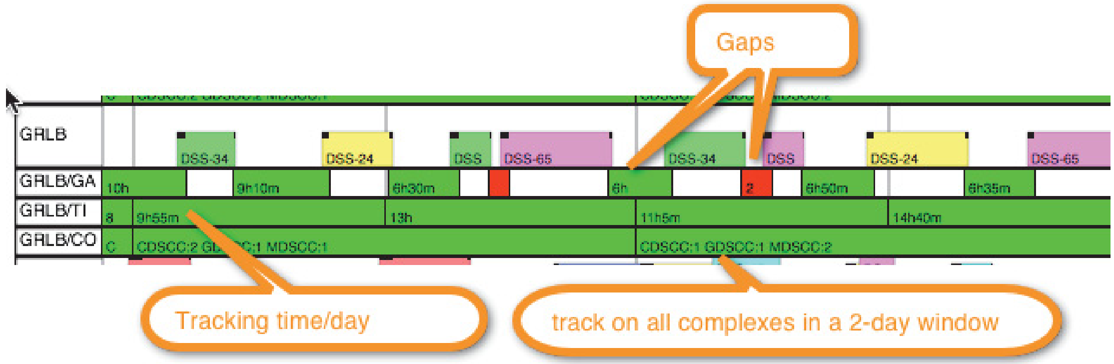

(a) Example of multiple timeline requirements applied to a single spacecraft, here GRAIL B, one of a pair of lunar orbiters. There is a gap constraint, a minimum tracking time constraint in a 24 hour UTC day, and a requirement to track on all three DSN complexes within a 48 hour period.

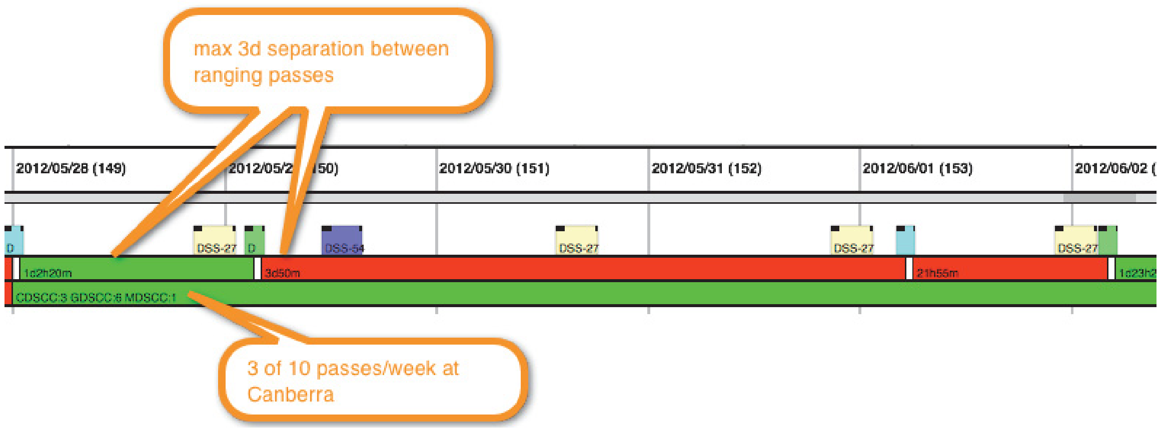

(b) Example of a gap constraint between ranging passes only, i.e. ignoring the intervening tracking passes. In this example, both the minimum and maximum gap requirement have been violated and the resulting intervals are colored red.

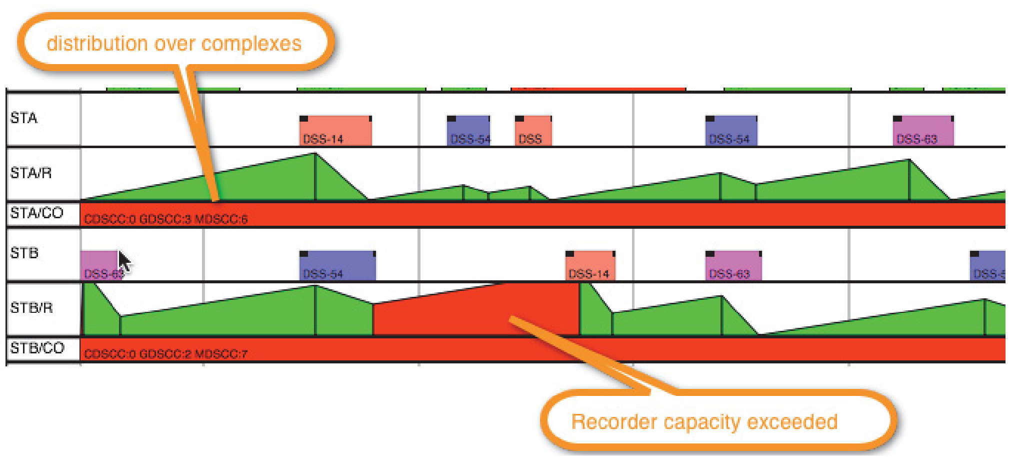

(c) Example of a recorder timeline constraint applied to the STEREO A/B mission pair, showing the violation of the constraint in one interval where the accumulated data would exceed the recorder capacity. Note that the recorder volume drops more quickly when a 70 meter contact is scheduled, due to the higher downlink data rate. The STERO spacecraft also have a requirement to schedule at least one track per week at each complex.

Figure 2. Timeline constraints for three representative mission sets, as depicted in the S^3 scheduling HTML5 GUI.

The Tracking Gaps timeline requirement applies to about a third of the DSN user set. In some cases, the gaps of concern are only for certain activity types, as illustrated in Figure 2b where gaps are only significant between subsequent ranging passes.

About 20% of users have DSN Complex Distribution requirements, but this varies depending on the phase of the mission. These requirements are typically driven by navigation considerations, where it is important to have ranging data from widely separated baselines in order to reduce ephemeris errors. Examples are shown in Figure 2a-c, where satisfaction or violation of the distribution requirement is clearly visible.

While most missions have onboard recorders, only a handful can potentially be modeled simply enough to include in the early stages of DSN scheduling. For those missions with uniform data collections rates and well-defined downlink rules, the Recorder timeline requirement can provide early visibility into recorder capacity and how it is affected by specific scheduling choices. An example is shown in Figure 2c for the STEREO A/B spacecraft.

By providing a highly visual view of these timeline constraints and preferences, users who are working on schedule changes to resolve conflicts can immediately see whether their proposed changes would introduce any violations. Presently, many scheduling users have custom scripts that they use to evaluate proposals from other users, but by providing for common models and visibility, feedback can be provided much more rapidly. This feedback has the potential to reduce the overall negotiation process effort and duration.

4. Long-Range Planning

While there are many similarities between the mid- and long-range planning and scheduling functions for DSN, there are also significant differences. Underlying both is the set of current and future DSN assets, including antennas and equipment, some coming into service and others being decommissioned. Both are based on DSN usage requirements from a varying mission set with a wide range of time-dependent tracking and navigation needs. Both are charged with arriving at an ultimately feasible allocation of DSN resources by balancing user needs and resolving periods of resource contention.

However, long-range planning also has some significant differences from mid-range:

- long-range planning has to deal with numerous and sometimes intrinsic sources of uncertainty, including:
  - unpredictable spacecraft locations for some missions and trajectory types, leading to uncertainties in visibility times from the different DSN antennas
  - unknown science targets beyond some time horizon in the future
  - uncertainties in the mission set, due to funding changes, launch date changes, or mission extensions
- optimization criteria and scenarios differ from mid-range, where the main objectives are to minimize conflicts in the schedule and violations of user requirements; for long-range planning a variety of other objectives may come into play, including:
  - identifying best times to schedule extended downtime for preventive maintenance, minimizing the impact on active missions
  - identifying best times to schedule special flexible but resource intensive operations, such as reference frame calibration activities
  - maximizing the satisfaction of requirements where, due to contention, not all requirements can be satisfied across the entire DSN user base

In addition, long range planning needs to provide information to mission planners about where contention with critical events may occur, so that this can be taken into account as early as possible in each mission's planning process. In many cases this needs to be provided during the mission proposal phase when, for both feasibility and costing, is is necessary to map out DSN allocation needs to some preliminary level of accuracy. Such proposal studies also impose a requirement for protection of proprietary or competition-sensitive information, whereas the midrange process for DSN allows general access to scheduling requirements and to the schedule itself.

Finally, long-range planning needs to support specification of a more abstract type of requirement with less detail than would be acceptable in mid-range. This serves two purposes: it represents at a coarse level some of the uncertainty in requirements, and it makes it easier to specify "what if" alternative scenarios.

The DSN has started on the first phase of a project to replace the current long-range planning tools with a new capability, designated *Loading Analysis and Planning Software* (LAPS), building on the functionality provided by S^3 for mid-range scheduling. LAPS will make direct use of a number of capabilities already deployed operationally in the mid-range S^3 software, including:

- the model of DSN asset availability including antennas and equipment, with time-varying availability for new construction or new types of equipment, and out-of-service dates for retired assets
- the model of DSN user and mission types, including
- ground- and space-based users, schedulable on non-interference basis or not
- multi-spacecraft constellations
- Multiple Spacecraft Per Antenna (MSPA) groupings and their special scheduling rules
- the service alias model, which defines what asset sets are allowable and preferable for a user, depending on the service desired
- the viewperiod model, specifying legal visibility intervals of various types, calculated by the Service Preparation System and imported in a form optimized for scheduling

- the scheduling requirement model, allowing (but not requiring) allocation needs to be specified to the same level of detail as mid-range requirements, should such detail be both available and necessary for the type of study to be undertaken
- the DSN Scheduling Engine algorithms used in the mid-range process, which would allow for fully detailed "what if" generation of hypothetical mid-range schedule periods in those cases where sufficient detail is available to warrant this level of analysis

Re-use of the S^3 software base in these areas provides a large degree of leverage in the development of LAPS, but several other areas are also being addressed with additional capabilities:

- a planning request representation to allow for more abstract and high-level specification of allocation needs than the scheduling requirement model allows (for example "3x 8hr tracks/week on 34m BWG for the 6 months of interplanetary cruise"); at the same time, planning requests will be convertible automatically into mid-range scheduling requests in order to minimize duplicate data entry and speed up the mid-range process
- the capability to define and run planning scenarios in an automated way, such as:
  - to assess a range of options for maintenance time
  - to evaluate nominal and fallback requirement options for resource contention periods
  - to quantify the impact of a mission's alternative launch dates on projected resource loading
- a multi-objective optimization mechanism to automatically generate a portfolio of candidate plans/schedules optimizing the tradeoffs among multiple quantitive objectives

The incorporation of multi-objective optimization into LAPS offers a new way to optimize DSN resource allocations, taking into account that there is no single objective that captures all of the disparate goals and objectives that are important. Multi-objective optimization has been employed in a wide variety of problem domains, including scheduling for science missions[13-16] and generating some requirements inputs to the DSN mid-range process[17].

The initial phase of LAPS development will encompass the modeling and optimization noted above. The second phase will extend the user interface elements of the software to allow end users, such as mission planners and schedulers, to directly enter their own planning require-

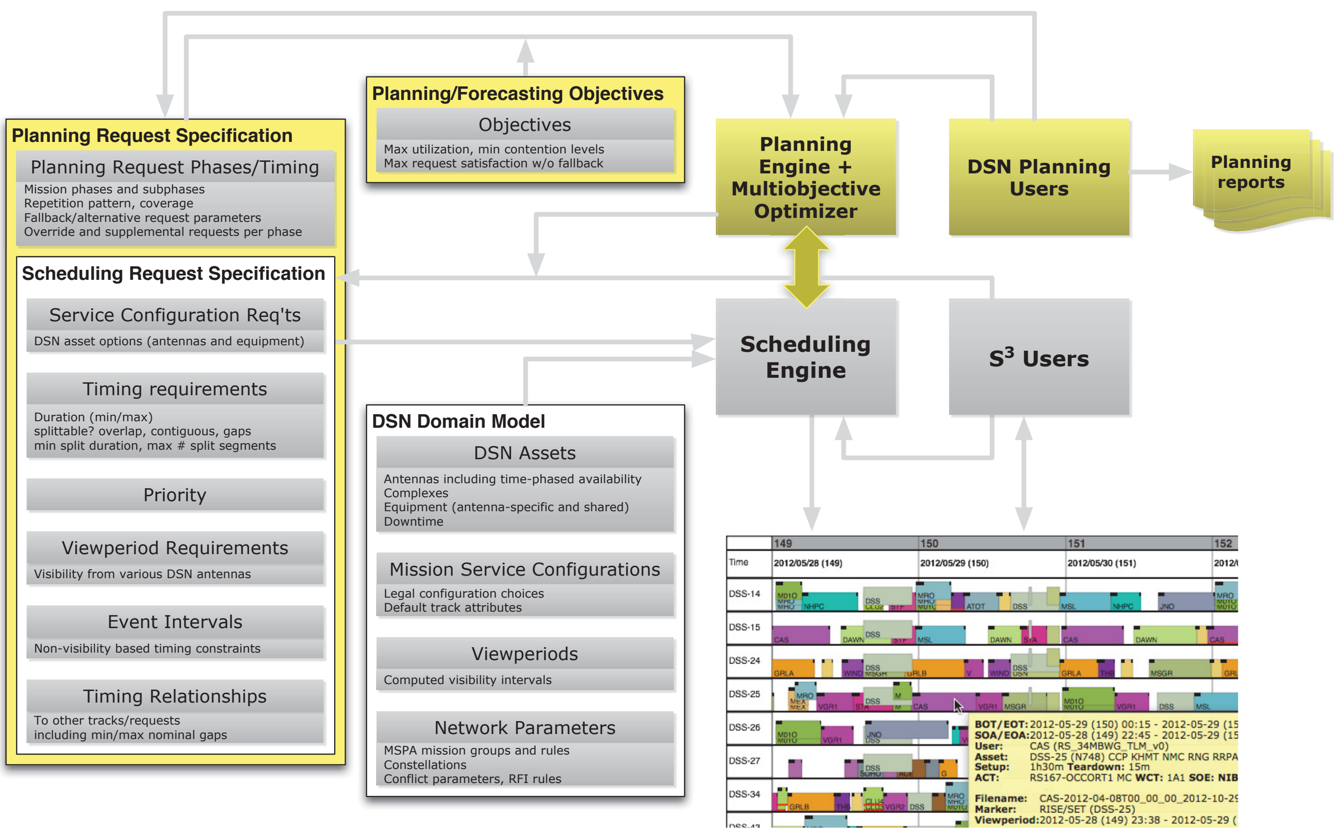

Figure 3. Extension of the S^3 data model to support long-range planning, forecasting, and downtime analysis. The existing S^3 data model and algorithms provide a basis that can be readily extended as indicated for long-range planning. Compare with Figure 1(b).

ments and conduct "what if" analyses using a baseline DSN asset and mission model. It will also include an extended report generation mechanism to generate a wider variety of tabular and graphical output formats.

5. Conclusions and Future Work

In this paper we have described the DSN scheduling process and software, including the initial operational deployment of the Service Scheduling Software (S^3) system, and its ongoing extension to support extended categories of scheduling requirements, as well as long-range planning and forecasting. S^3 represents a new approach to scheduling the DSN, embodying a request-driven approach to scheduling along with a collaborative peer-to-peer negotiation environment using modern web application and database technology. Future work is expected to address a number of areas including:

- extension to real-time scheduling – this third phase of the DSN scheduling process covers the period from execution out to some small number of weeks in the future. Extending S^3 to support this phase involves some challenging technical problems of integration with existing systems and support for contingency scheduling (e.g. launch slips, unplanned asset downtime); at the same time, bringing the information model of S^3 into the real-time domain will allow for decision making considering options that are not now accessible
- cross-network scheduling – NASA has recommended[18] integrating access to the capabilities provided by its three major networks: DSN, the Space Network (SN), and the Near Earth Network (NEN). For those users requiring services from two or all three of these networks, such integration would be a source of significantly improved efficiency and cost savings. S^3 has the potential to serve as a common scheduling platform in this regard.

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. We gratefully acknowledge the support of the DSN scheduling community over the course of this work.

Bibliography

1. Berner, J.B. and J. Statman, *Increasing the cost-efficiency of the DSN*, in *SpaceOps 2008*2008: Heidelberg, Germany.
2. Imbriale, W.A., *Large Antennas of the Deep Space Network*2003: Wiley.
3. Lacey, N. and D.G. Morris, *JPL RAPSO Long-Range Forecasting*, in *12th AAS/AIAA Space Flight Mechanics Meeting*2002: San Antonio, Texas.
4. Morris, D. and E. Burke, *The JPL Resource Allocation Planning and Scheduling Office (RAPSO) process*, in *12th AAS/AIAA Space Flight Mechanics Meeting*2002, AAS: San Antonio, Texas.
5. Johnston, M.D., et al., *Request-Driven Schedule Automation for the Deep Space Network*, in *SpaceOps 2010*2010, AIAA: Huntsville, AL.
6. Johnston, M.D. and D. Tran, *Automated Scheduling for NASA's Deep Space Network*, in *6th International Workshop on Planning and Scheduling in Space (IWPSS)*2011: Darmstadt, Germany.
7. Johnston, M.D., et al., *Automating Mid- and Long-Range Scheduling for NASA's Deep Space Network*, in *SpaceOps 2012*2012: Stockholm, Sweden.
8. Clement, B.J. and M.D. Johnston. *Design of a Deep Space Network Scheduling Application*. in *International Workshop on Planning and Scheduling for Space (IWPSS-06)*. 2006. Baltimore, MD: Space Telescope Science Institute.
9. Johnston, M.D. and B.J. Clement. *Automating Deep Space Network Scheduling and Conflict Resolution*. in *ISAIRAS-05*. 2005. Munich, Germany.
10. Carruth, J., et al., *A Collaborative Scheduling Environment for NASA's Deep Space Network*, in *SpaceOps 2010*2010, AIAA: Huntsville, AL.
11. Johnston, M.D., et al. *Request-Driven Scheduling for NASA's Deep Space Network*. in *International Workshop on Planning and Scheduling for Space (IWPSS)*. 2009. Pasadena, CA.
12. Chien, S., et al. *ASPEN - Automating Space Mission Operations using Automated Planning and Scheduling*. in *SpaceOps 2000*. 2000. Toulouse, France.
13. Johnston, M.D. *Multi-Objective Scheduling for NASA's Deep Space Network Array*. in *International Workshop on Planning and Scheduling for Space (IWPSS-06)*. 2006. Baltimore, MD: Space Telescope Science Institute.
14. Johnston, M.D., *An Evolutionary Algorithm Approach to Multi-Objective Scheduling of Space Network Communications.* International Journal of Intelligent Automation and Soft Computing, 2008. **14**: p. 367-376.
15. Johnston, M.D. and M. Giuliano. *MUSE: The Multi-User Scheduling Environment for Multi-Objective Scheduling of Space Science Missions*. in *IJCAI Workshop on Space Applications of AI*. 2009. Pasadena, CA.
16. Giuliano, M. and M.D. Johnston. *Multi-Objective Evolutionary Algorithms for Scheduling the James Webb Space Telescope*. in *International Conference on Automated Planning and Scheduling (ICAPS)*. 2008. Sydney, Australia.
17. Johnston, M.D. and M. Giuliano. *Multi-Objective Scheduling for the Cluster II Constellation*. in *6th International Workshop on Planning and Scheduling in Space (IWPSS)*. 2011. Darmstadt, Germany.
18. SCAWG, *NASA Space Communications and Navigation Architecture Recommendations, Final Report*, 2006, Space Communications Architecture Working Group (SCAWG).